

The Bundled Crossing Number

Md. Jawaherul Alam¹, Martin Fink², and Sergey Pupyrev^{3,4}

¹ Department of Computer Science, University of California, Irvine

² Department of Computer Science, University of California, Santa Barbara

³ Department of Computer Science, University of Arizona, Tucson

⁴ Institute of Mathematics and Computer Science, Ural Federal University

Abstract. We study the algorithmic aspect of edge bundling. A bundled crossing in a drawing of a graph is a group of crossings between two sets of parallel edges. The bundled crossing number is the minimum number of bundled crossings that group all crossings in a drawing of the graph. We show that the bundled crossing number is closely related to the orientable genus of the graph. If multiple crossings and self-intersections of edges are allowed, the two values are identical; otherwise, the bundled crossing number can be higher than the genus.

We then investigate the problem of minimizing the number of bundled crossings. For circular graph layouts with a fixed order of vertices, we present a constant-factor approximation algorithm. When the circular order is not prescribed, we get a $\frac{6c}{c-2}$ -approximation for a graph with n vertices having at least cn edges for $c > 2$. For general graph layouts, we develop an algorithm with an approximation factor of $\frac{6c}{c-3}$ for graphs with at least cn edges for $c > 3$.

1 Introduction

For many real-world networks with substantial numbers of links between objects, traditional graph drawing algorithms produce visually cluttered and confusing drawings. Reducing the number of edge crossings is one way to improve the quality of the drawings. However, minimizing the number of crossings is very difficult [5, 7], and a large number of crossings is sometimes unavoidable. Another way to alleviate this problem is to employ the edge bundling technique in which some edge segments running close to each other are collapsed into bundles to reduce the clutter [8, 12, 16, 19, 20, 21, 25]. While these methods produce simplified drawings of graphs and significantly reduce visual clutter, they are typically heuristics and provide no guarantee on the quality of the result.

We study the algorithmic aspect of edge bundling, which is listed as one of the open questions in a recent survey on crossing minimization by Schaefer [26]. Our goal is to formalize the underlying geometric problem and design efficient algorithms with provable theoretical guarantees. In our model, *pairwise* edge crossings are merged into bundles of crossings, reducing the number of *bundled crossings*, where a bundled crossing is the intersection of two groups of edges; see Fig. 1. We consider both the general setting, where multiple crossings and self-intersections of the edges are allowed, and the more natural restricted setting in which only simple drawings are allowed.




Fig. 1: Circular layout of the Chvátal graph: (a) 28 pairwise edge crossings, (b) 13 bundled crossings.

1.1 Our Contribution

We first prove that in the most general setting (when a pair of edges is allowed to cross multiple times and an edge may be crossed by itself or by an incident edge) the bundled crossing number coincides with the orientable genus of the graph (Section 2); thus, computing it exactly is NP-hard [29]. In the more natural setting restricted to simple drawings—without double- and self-crossings—, the bundled crossing number of some graphs is strictly greater than the genus.

Next, we consider the *circular bundled crossing number* (Section 3), that is, the minimum number of bundled crossings that can be achieved in a circular graph layout. For a fixed circular order of vertices, we present a 16-approximation algorithm and a fixed-parameter algorithm with respect to the number of bundled crossings. For circular layouts without a given vertex order, we develop an algorithm with the approximation factor $\frac{6c}{c-2}$ for graphs with n vertices having at least cn edges for $c > 2$.

In Section 4, we study the *bundled crossing number* for general drawings. The algorithm for circular layouts can also be applied for this setting; we show that it guarantees the approximation factor $\frac{6c}{c-3}$ for graphs with at least cn edges for $c > 3$. We then suggest an alternative algorithm that produces fewer bundled crossings for graphs with a large planar subgraph.

Finally, by extending our analysis for circular layouts, we resolve one of the open problems stated by Fink et al. [15] for an ordering problem of paths on a graph arising in visualizing metro maps (Section 5).

1.2 Related Work

Edge crossings. Crossing minimization is a rich topic in graph drawing [5] but still poorly understood from the algorithmic point of view. The best currently known algorithm implies an $\mathcal{O}(n^{9/10})$ -approximation for the minimum crossing number on graphs having bounded maximum degree [7]. In contrast, the problem is NP-hard even for cubic graphs and a hardness of constant-factor approximation is known [6]. Minimizing crossings in circular layouts is also NP-hard, and

several heuristics have been proposed [3, 17]. For graphs with $m \geq 4n$, an $\mathcal{O}(\log^2 n)$ -approximation algorithm exists [28]. Our algorithm guarantees an $\mathcal{O}(1)$ -approximation for bundled crossings under that condition.

Bundled crossings are closely related to the model of *degenerate crossings* in which multiple edge crossings at the same point in the plane are counted as a single crossing if all pairs of edges passing through the point intersect. An unrestricted variant, called the *genus crossing number* ($\text{gcr}(G)$), allows for self-crossings of edges and multiple crossings between pairs of edges. Mohar showed that the genus crossing number equals the *non-orientable genus* of a graph [23]; thus, $\text{gcr}(G) = \mathcal{O}(m)$. This is similar to our result that the bundled crossing number in this unrestricted setting equals the *orientable genus* of the graph. If self-crossings are not allowed, then we obtain the *degenerate crossing number* ($\text{dcr}(G)$) [24, 27]. It was conjectured by Mohar [23] that the genus crossing number always equals the degenerate crossing number; Schaefer and Štefankovič show that $\text{dcr}(G) \leq 6 \cdot \text{gcr}(G) = \mathcal{O}(m)$. A further restriction of the problem forbids multiple crossings between a pair of edges. The corresponding *simple degenerate crossing number* is $\Omega(m^3/n^2)$ for graphs with $m \geq 4n$ edges [1]. Thus multiple crossings between pairs of edges are significant for the corresponding value of the crossing number. Notice the difference to the bundled crossing number, which is always $\mathcal{O}(m)$, even when no self- and multiple crossings are allowed.

Recently, Fink et al. [14] introduced the bundled crossing number. However, they only study the bundled crossing number of a given embedding and show that determining the number is NP-hard. They also present a heuristic that in some cases, e.g., in circular layouts, yields a constant-factor approximation. In contrast, we study the variable-embedding setting: minimize the bundled crossing number over all embeddings of a graph, which is posed as an open problem in [14].

Edge bundling. Improving the quality of layouts via edge bundling is related to the idea of confluent drawings, when a non-planar graph is presented in a planar way by merging groups of edges [10, 11]. The first discussion of bundled edges in the graph drawing literature appeared in [17], where the authors improve circular layouts by routing edges either on the outer or on the inner face of a circle. The hierarchical approach by Holten [19] bundles the edges based on an additional tree structure, and the method is also applied for circular layouts. Similar to [11, 17, 19], we study circular graph layouts. Edge bundling methods for general graph layouts are suggested in [8, 12, 16, 20, 21]. While these methods create an overview drawing, they allow the edges within a bundle to cross and overlap each other arbitrarily, making individual edges hard to follow. The issue is addressed in [4, 25], where the edges within a bundle are drawn parallel, as lines in metro maps. To the best of our knowledge, none of the above works on edge bundling provides a guarantee on the quality of the result, though they can be applied in conjunction with our algorithms to provide a better visualization.

Metro maps. Crossing minimization has also been studied in the context of visualizing metro maps. There, a planar graph (the metro network) and a set of paths in the graph (metro lines) are given. The goal is to order the paths

along the edges of the graph so as to minimize the number of crossings. Fink, Pupyrev, and Wolff [15] suggest to merge single line crossings into crossings of blocks of lines minimizing the number of *block crossings* in the map. They devise approximation algorithms for several classes of simple underlying networks (paths, upward trees) and an asymptotically worst-case optimal algorithm for general networks. While we use some ideas of [15] (Section 3.1), bundled crossings are more general, since the edges are not restricted to be routed along a specified planar graph. Furthermore, we resolve an open question stated in [15].

2 Bundled Crossings and Graph Genus

Let $G = (V, E)$ with $n = |V|$ and $m = |E|$ be a graph drawn in the plane (with crossings). A *bundled crossing* is a subset C of the crossings so that the following conditions hold:

- (i) Every crossing in C belongs to edges $e_1 \in E_1$ and $e_2 \in E_2$, for two subsets $E_1, E_2 \subseteq E$ (E_1 and E_2 are the *bundles* of the bundled crossing), and C contains a crossing of each edge pair e_1, e_2 , for $e_1 \in E_1$ and $e_2 \in E_2$.
- (ii) One can find a pseudodisk D —a closed polygonal region crossing every edge at most twice—that separates C (in its interior) from all remaining crossings of the embedding. No edge $e \notin E_1 \cup E_2$ intersects D . The requirement ensures that the bundled crossing is visually separated from the rest of the drawing.

The bundled crossing number of a drawing is the minimum number of bundled crossings into which the crossings can be partitioned (with disjoint pseudodisks). The *bundled crossing number* $bc(G)$ of G is the minimum number of bundled crossings in a drawing of G . For a circular layout, we denote the *circular bundled crossing number* by $bc^\circ(G)$. If the circular order π of vertices is prescribed, we speak of the *fixed circular bundled crossing number*, $bc^\circ(G, \pi)$. Clearly, $bc(G) \leq bc^\circ(G) \leq bc^\circ(G, \pi)$.

We now discuss the relation of the bundled crossing number to the orientable genus of the graph. More specifically, consider the unrestricted drawing style for graphs in which double crossings of edges are allowed, as well as self intersections and crossings of adjacent edges. Let $bc'(G)$ be the minimum number of bundled crossings achievable for G in this unrestricted drawing style. We show that $bc'(G)$ equals the graph genus.

Theorem 1. *For every graph G with genus $g(G)$, it holds that $bc'(G) = g(G)$.*

Proof. It is easy to show that $g(G) \leq bc'(G)$. We take a drawing of G with the minimum number of bundled crossings, $bc'(G)$, on the sphere. Then, for every bundled crossing, we add a handle to the sphere, where we route one of the bundles through the handle and one on top of it. This way we get a crossing-free drawing of G on a surface of genus $bc'(G)$.

For the other direction, assume that we have a crossing-free drawing of G on a surface of genus $g = g(G)$. It is known that such a drawing can be modeled using the representation of a genus- g surface by a *fundamental polygon* with $4g$ sides in the plane [22]. More precisely, the sides of the polygon are numbered in circular




Fig. 2: K_6 drawn in a fundamental square modeling a torus.




Fig. 3: A single bundled crossing outside the fundamental polygon.

order $a_1, b_1, a'_1, b'_1, \dots, a_g, b_g, a'_g, b'_g$; for $1 \leq k \leq g$, the pairs (a_k, a'_k) and (b_k, b'_k) of sides are identified in opposite direction, meaning that an edge leaving side a_k appears on the corresponding position of edge a'_k ; see Fig. 2 for an example showing K_6 drawn in a fundamental square that models a drawing on the torus. Directly transforming a drawing on the surface into the fundamental polygon can lead to vertices appearing multiple times on the polygon's boundary; however, small movements of the vertices on the surface fix this. Thus, we assume that all vertices lie in the interior of the fundamental polygon, and all edges leave the polygon only in the relative interior of a side of the polygon; especially, every point of an edge appears at most twice on the boundary of the fundamental polygon. (There can be parts of edges connecting two points on different sides of the polygon without directly touching a vertex as in Fig. 2).

Given such a crossing-free representation of the drawing of G via the fundamental polygon, we create a new drawing of G in the plane by connecting parts of the edges outside of the fundamental polygon. For every $1 \leq k \leq g$, we connect identified points of edges on a_k, a'_k, b_k , and b'_k as shown in Fig. 3. It is easy to see that for every k , only one bundled crossing is necessary; furthermore, all g tuples of four consecutive sides are independent. Hence, we get a drawing with g bundled crossing, which proves that $\text{bc}'(G) \leq g(G)$. \square

When creating a drawing as in the second part of the above proof, it may happen that we introduce (i) double crossings of edges, (ii) crossings between adjacent edges, or (iii) self intersections of an edge. Certainly, a drawing avoiding such configurations—that is, a *simple* drawing—is preferred. From now on, we only consider simple drawings. Let $\text{bc}(G)$ denote the minimum number of bundled crossings achievable with a simple drawing of G . It turns out that insisting on a simple drawing sometimes makes additional bundled crossing necessary.

Lemma 1. *For every graph $G = (V, E)$, $\text{bc}(G) \geq g(G)$, and there are graphs G for which $\text{bc}(G) > g(G)$.*

Proof. Since we only restrict the allowed drawings, we clearly have $\text{bc}(G) \geq \text{bc}'(G) = g(G)$ and the first claim follows.

For the second part of the lemma, consider the complete graph on six vertices, K_6 , with genus $g(K_6) = 1$; there is a crossing-free drawing of K_6 on the torus. Every realization of K_6 with only one bundled crossing leads to a drawing on the torus. Consider such a drawing in the fundamental polygon model of the torus; in this case, the fundamental polygon can be seen as an axis-aligned square where edges can go to the upper, lower, left, and right side of the square. If two edges incident to the same vertex v leave the square to adjacent sides, the edges cross in the bundled crossing, which is forbidden. Furthermore, no part of an edge can enter and leave the square on adjacent sides since this would result in a forbidden self-intersection. Given these constraints, it is not hard but technical to verify that K_6 cannot be embedded on the torus and, therefore, $\text{bc}(K_6) > 1$. We refer to Lemma 7 in Appendix A. \square

It is easy to see that $g(G) = \mathcal{O}(m)$ by introducing a handle on the sphere for each edge. Furthermore, for the complete graph K_n , it is known that $g(K_n) = \lceil (n-3)(n-4)/12 \rceil$, that is, $g(G) = \Theta(m)$ for some graphs. Clearly, we cannot do better with bundled crossings, that is, $\text{bc}(G) = \Omega(m)$ for some graphs. In Section 3.1 we show that $\mathcal{O}(m)$ bundled crossings always suffice, even if we are using a circular layout with a fixed order of vertices. This means that for complete graphs, all bundled crossing number variants and the genus are within a constant factor from each other. An interesting question is how large the ratio between the bundled crossing number and the graph genus can get for general graphs.

It is known that $\Omega(m^3/n^2)$ single crossings are necessary for graphs with n vertices and $m \geq 4n$ edges [2]. For dense graphs with $m = \Theta(n^2)$ edges, $\Theta(m^2)$ crossings are required, while the bundled crossing number is $\mathcal{O}(m)$. Therefore, using edge bundles can significantly reduce visual complexity of a drawing.

3 Circular Layouts

Now we consider circular graph layouts. Let $G = (V, E)$ be a graph and let $\pi = [v_1, \dots, v_n]$ be a permutation of its vertices. The goal is to draw G in such a way that the vertices are placed on the boundary of a disk in the circular order prescribed by π , all edges are drawn inside the circle, and the number of bundled crossings, $\text{bc}^\circ(G, \pi)$, is minimized. We start with a scenario when π is predefined.

3.1 Circular Layouts with Fixed Order

Since in our model adjacent edges are not allowed to cross and the circular order of the vertices is fixed, the order of outgoing edges for every vertex is unique for π . Hence, we may assume that G is a matching. Note that in this case the circular layout can be seen as a *weak pseudoline arrangement*, that is, an arrangement of pseudolines in which not every pair of pseudolines has to cross [9].

Assume that edges e_1 and e_2 are *parallel*, that is, they do not have to cross, and they start and end as immediate neighbors. Clearly, in any simple drawing,

e_1 and e_2 do not cross and they are crossed by exactly the same set of other edges; otherwise we would have a forbidden double crossing. Therefore, we can remove e_2 from the instance, find a drawing for the remaining graph, and then reintroduce e_2 without an additional bundled crossing. To this end, we route e_2 parallel to e_1 and let it participate in e_1 's bundled crossings in the same bundle as e_1 . Thus, we may assume that (i) the input contains no parallel pairs of edges. Additionally, we assume that (ii) every edge of the input graph has to be crossed by an edge (which can be checked by looking at the given circular order); otherwise, such an edge is removed from the input and later reinserted without crossings. In the following we assume that the input satisfied both conditions (i) and (ii) and such a graph is called *simplified*.

Next we develop an approximation algorithm for $\text{bc}^\circ(G, \pi)$ by showing how to find a solution with only a linear number of bundled crossings, and proving that every feasible solution, even an optimum one, must have a linear number of bundled crossings. We start with the lower bound.

Lemma 2. *Let $G = (V, E)$ be a simplified graph with fixed circular vertex order π . Then, $\text{bc}^\circ(G, \pi) \geq m/16$.*

Proof. Assume we are given a circular drawing of G with the minimum number of bundled crossings. Such a drawing is a weak pseudoline arrangement. Let H be the embedded planar graph that we get by planarizing the drawing, that is, by replacing each crossing by a crossing vertex and adding the cycle (v_1, v_2, \dots, v_n) . We consider the faces of H . Some faces are bounded by original edges and an additional edge stemming from the cycle. Next we lower bound the number of triangles in the pseudoline arrangement and, hence, the triangular faces in H .

Assume that we follow some edge in the drawing and analyze the faces at one of its sides. If all faces were quadrilaterals, then the edge would be completely parallel to a neighboring edge, which is not possible in a simplified instance. Hence, on both sides of the edge we find at least one face that is either a triangle or a k -gon with $k \geq 5$. For $k \geq 3$, let f_k be the number of faces in the drawing of H of degree k . Since we see at least $2m$ sides of such faces and every side only once, we have $2m \leq 3f_3 + \sum_{k \geq 5} kf_k$. Fink et al. [14] show that $f_3 = 4 + \sum_{k \geq 5} (k-4)f_k$. Hence, $2m \leq 3f_3 + \sum_{k \geq 5} (k-4)f_k + 4 \sum_{k \geq 5} f_k \leq 3f_3 + (f_3 - 4) + 4(f_3 - 4) \leq 8f_3$, which implies $f_3 \geq m/4$. Note that the bound is tight; see Appendix B.

To complete the proof of the lemma, we use a result of Fink et al. [14], who show that the crossings in a fixed drawing can be partitioned into no less than $f_3/4$ bundled crossings. Since every drawing has at least $m/4$ triangles, $\text{bc}^\circ(G, \pi) \geq m/16$. \square

Note that, as Fink et al. [14] point out, there exist circular drawings whose crossings can be partitioned into no less than $\Theta(m^2)$ bundled crossings. However, we can choose the drawing as long as we follow the cyclic order, π , of vertices. We use this freedom and show how to construct a solution with $\mathcal{O}(m)$ bundled crossings.

Lemma 3. *Let $G = (V, E)$ be a graph with a fixed circular vertex order. We can find a circular layout with at most $m - 1$ bundled crossings in $\mathcal{O}(m^2)$ time.*




Fig. 4: Finding a circular layout with $m - 1$ bundled crossings (gray shaded).

Proof. Recall that we may assume that the input graph is a matching. Since only the circular order of the vertices matters for the combinatorial embedding, we transform the circle into a rectangle with v_1, \dots, v_n placed on the lower side from left to right; see Fig. 4. We produce a drawing in which every edge $e = (v_i, v_j)$ with $i < j$ consists of two straight-line segments.⁵ The first segment leaves v_i with a slope α ; when the segment is above v_j it is followed by a vertical segment connecting down to v_j . Since there are only two slopes, every crossing is between a vertical segment and a segment of slope α . It is easy to see that two edges (v_i, v_j) and $(v_{i'}, v_{j'})$ cross if the endvertices are interleaved, that is, if $i < i' < j < j'$ or $i' < i < j' < j$. In that case, the edges have to cross in any possible embedding and we do not introduce additional crossings.

Finally, we create a single bundled crossing for each edge e consisting of all crossings of e 's vertical segment. It is easy to see that this yields a feasible partitioning of all crossings into bundled crossings. Since the edge ending at vertex v_n will not have any crossing on its vertical segment, the number of bundled crossings is at most $m - 1$. The drawing is created in $\mathcal{O}(m)$ time but the time needed to produce a combinatorial embedding depends on the number of crossings; it is bounded by $\mathcal{O}(m^2)$. \square

The upper bound of $m - 1$ is tight: a matching in which every edge crosses every other edge requires that many bundled crossings. Combining the algorithm and the lower bound of Lemma 2, we get the following result.

Theorem 2. *For a graph G with a fixed circular vertex order, we can find a 16-approximation for the fixed circular bundled crossing number in $\mathcal{O}(m^2)$ time.*

Fixed-Parameter Tractability. We now show that deciding whether a solution with at most k bundled crossings exists is fixed-parameter tractable with respect to k . The crucial instruments for achieving this are the graph simplification and the lower bound of Lemma 2. If after the simplification, G has more than $16k$ edges, we know that $\text{bc}^\circ(G, \pi) > 16k/16 = k$ and we can reject the instance. Otherwise, if at most k edges remain, we can afford to solve the problem exhaustively.

Theorem 3. *Let $G = (V, E)$ be a graph with a fixed circular vertex order π . Deciding whether $\text{bc}^\circ(G, \pi) \leq k$ is fixed-parameter tractable with respect to k with a running time of $\mathcal{O}(2^{0.657k^2} k^{128k^2} + m)$.*

⁵ We thank an anonymous reviewer for suggesting this simplified proof.

Proof. We simplify the graph in $\mathcal{O}(m)$ time. Afterwards, we check every combination of circular order, combinatorial embedding, and partitioning of the crossings into up to k sets. If any such combination yields a feasible partitioning into bundled crossings, we accept the instance; otherwise, we reject it.

There are at most $\binom{16k}{2} \leq 128k^2$ pairs of edges that need to cross. Hence, there are up to k^{128k^2} ways to partition the crossings into up to k sets. Since every pair of edges crosses at most once, the circular embedding can be extended to a pseudoline arrangement (in which every pair crosses exactly once). Felsner and Valtr proved [13] that there are at most $2^{0.657k^2}$ arrangements of k pseudolines, and Yamanaka et al. [30] presented a method that iterates over all pseudoline arrangements using $\mathcal{O}(k^2)$ total space and $\mathcal{O}(1)$ time per arrangement. For each pseudoline arrangement, we can check whether an embedding with the prescribed circular order occurs as a part in $\mathcal{O}(k^3)$ time; within the same time bound, we can check whether a given partitioning of the crossings yields feasible bundled crossings. In total this takes $\mathcal{O}(2^{0.657k^2} k^{128k^2} + m)$ time. \square

3.2 Circular Layouts with Free Order

We now study the variant of the problem in which the circular order of the vertices is not known. How can one find a suitable order? A possible approach would be finding an order that optimizes some aesthetic criteria (e.g., the total length of the edges [17] or the number of pairwise crossings [3]) and then applying the algorithm of Lemma 3. Next we analyze such an approach.

In Section 2, we have already seen that $\text{bc}(G) \geq \text{g}(G)$. We can use this for getting a lower bound for the bundled crossing number.

Lemma 4. *For every graph $G = (V, E)$ with n vertices and m edges, $\text{bc}(G) \geq (m - (3n - 6)) / 6$ and $\text{bc}^\circ(G) \geq (m - (2n - 3)) / 6$.*

Proof. Assume we have a crossing-free drawing of graph G on a surface of genus $g = \text{g}(G)$. The relation between vertices, edges, and faces is described by the Euler formula $n - m + f = 2 - 2g$. Combining this with $2m \geq 3f$, we get that $\text{bc}(G) \geq \text{g}(G) \geq (m - (3n - 6)) / 6$.

Now consider a circular drawing with the minimum number $k = \text{bc}^\circ(G)$ of bundled crossings. All n vertices lie on the outer face. Hence, we can add $n - 3$ edges triangulating the outer face without introducing new crossings. We get a new graph G' with $m' = m + n - 3$ edges and a (non-circular) drawing of G' with k bundled crossings. Hence, $k \geq (m' - (3n - 6)) / 6 = (m - (2n - 3)) / 6$. \square

For dense graphs with more than $2n$ edges, we can get a constant-factor approximation using the upper bound of $m - 1$ with an arbitrary order.

Theorem 4. *Let $G = (V, E)$ be a graph with $m \geq cn$ for some $c > 2$. There is an $\mathcal{O}(n^2)$ -time algorithm that computes a solution for the circular bundled crossing number with an approximation factor of $\frac{6c}{c-2}$.*

Proof. Using the algorithm of Lemma 3, we find a solution with at most $m - 1$ bundled crossings. By Lemma 4, $(m - (2n - 3))/6$ crossings are required. Then the approximation factor is $\frac{m-1}{(m-(2n-3))/6} = 6 \left(1 + \frac{2n-4}{m-2n+3}\right) \leq 6 \left(1 + \frac{2n}{m-2n}\right) \leq 6 \left(1 + \frac{2}{c-2}\right) = \frac{6c}{c-2}$, which is constant for every $c > 2$ and $n \geq 1$. \square

For constructing a constant-factor approximation algorithm for sparse graphs with $m \leq 2n$ one would need better bounds. We next suggest a possible direction by improving our algorithm for some input graphs. The idea is to save some crossings by first drawing an outerplanar subgraph of G .

Lemma 5. *Let $G = (V, E)$ be a graph and $G^* = (V, E^*)$ be a subgraph of G having $m^* = |E^*|$ edges that is outerplanar with respect to a vertex order π . Then $\text{bc}^\circ(G) \leq \text{bc}^\circ(G, \pi) \leq 2(m - m^*)$ and we can find such a solution in $\mathcal{O}(m^2)$ time.*

Proof. The algorithm is similar to the one used in Lemma 3 in which every edge consists of two segments. This time we initialize the embedding by adding the edges of E^* without crossings, each with a segment of slope α . Next, we add the remaining edges from left to right ordered by their first vertex. When adding edge $e = (v_i, v_j)$ with $i < j$, we route the edge with two vertical segments and a middle segment of slope α . We start upward from v_i so that the first segment crosses all edges present at $x = x(v_i)$ that have to cross e , but no other edge. We start the middle segment with slope α there and complete with a vertical segment at $x = x(v_j)$. It is easy to see that any edge of E^* whose vertical segment could intersect e must start left of v_i . However, our routing of e places the possible crossing on a vertical segment of e . Hence, all vertical segments of edges of E^* are crossing-free. Creating a bundled crossing for each vertical segment of the edges of $E - E^*$ results, therefore, in at most $2(m - m^*)$ bundled crossings. \square

This bound is asymptotically tight; see Appendix D.

4 General Drawings

We now consider general (non-circular) drawings. Note that Lemma 4 provides a lower bound for the bundled crossing number, and Lemma 3 gives an algorithm that can be applied for general drawings. Combining the lower and the upper bounds, we get the following result for dense graphs.

Theorem 5. *Let $G = (V, E)$ be a graph with $m \geq cn$ for some $c > 3$. There is an $\mathcal{O}(n^2)$ -time algorithm that computes a solution for the bundled crossing number with an approximation factor of $\frac{6c}{c-3}$.*

Proof. By Lemma 4, $\text{bc}(G) \geq (m - (3n - 6))/6$, and by Lemma 3, $\text{bc}(G) \leq m - 1$. Then the approximation factor of the algorithm of Lemma 3 is $\frac{m-1}{(m-(3n-6))/6} = 6 \left(1 + \frac{3n-7}{m-3n+6}\right) \leq 6 \left(1 + \frac{3n}{m-3n}\right) \leq 6 \left(1 + \frac{3}{c-3}\right) = \frac{6c}{c-3}$. \square




Fig. 5: Adding edge $e = (u, v)$ requires two circular arcs.




Fig. 6: Inserting edge $e = (u, v)$ with two circular arcs.

Can we improve the algorithm for general drawings? Next we develop an alternative upper bound based on a planar subgraph $G^* = (V, E^*)$ of G , which produces fewer bundled crossings if $m^* = |E^*| > 3m/4$.

Lemma 6. *Let $G = (V, E)$ be a graph, let $G^* = (V, E^*)$ be its planar subgraph, and let $m^* = |E^*|$. Then, $\text{bc}(G) \leq 4(m - m^*)$.*

Proof. We start with a topological book embedding of G^* , that is, a planar embedding with all vertices on the x -axis and the edges composed of circular arcs whose center is on the x -axis. Giordano et al. [18] show how to construct such an embedding with at most two circular arcs per edge and all edges being x -monotone (that is, edges with two circular arcs cannot change the direction).

We add the edges of $E' = E \setminus E^*$ to get a non-planar topological book embedding (with up to two circular arcs per edge) and keep the drawing simple, that is, free of self-intersections, double crossings, and crossings of adjacent edges. Then we split the drawing at the spine and interpret each half as a circular layout with fixed order. Using the algorithm of Lemma 5, we get an embedding with at most $2(m - m^*)$ crossings for each side and $4(m - m^*)$ crossings in total.

It remains to show how to add an edge $e = (u, v) \in E'$. Consider all planar edges incident to u and v . If we can add e as a single circular arc above or below the spine without crossing any of these edges, we do so. Otherwise, two edges e_1 adjacent to u and e_2 adjacent to v exist (see Fig. 5), and e must be inserted using two circular arcs. We consider all these obstructing two-bend edges incident to u and v and insert e by placing its bend next to the rightmost bend of an edge incident to u (see Fig. 6), avoiding all intersections with planar edges. Bends of the edges incident to u are ordered by their endvertex so that they do not cross.

It is easy to see that there are no self-intersections and no crossings of adjacent edges. There are also no double crossings: Otherwise, let e_1 and e_2 be a pair of edges that cross both above and below the spine. Assume that $e_1 \in E'$, $e_2 \in E^*$. Since e_1 consists of two segments, there must be adjacent planar edges that caused e_1 's shape. We find such an edge e'_1 that crosses with the planar edge e_2 , a contradiction; see Fig. 7a. If $e_1, e_2 \in E'$, we find a planar edge e'_2 causing the two-arc shape of e_2 , such that e'_1 and e'_2 cross, another contradiction; see Fig. 7b.

□




Fig. 7: Double crossings of edges are not possible

5 Block Crossings in Metro Maps

Our analysis has an interesting application for block crossings in metro maps [15]. The block crossing minimization problem (BCM) asks to order simple paths (metro lines) along the edges of a plane graph (underlying metro network) so as to minimize the total number of block crossings. Fink et al. [15] present a method that uses two block crossings per line on a tree network, and ask whether a (constant-factor) approximation is possible. With the help of the lower bound of Lemma 2, we affirmatively answer the question. We provide a sketch of the proof; see Appendix C for details.

Theorem 6. *There is an $\mathcal{O}(\ell^2)$ -time 32-approximation algorithm for BCM, where ℓ is the number of metro lines and the underlying network is a tree.*

Proof. Suppose that we have a solution with k block crossings on the tree. We can interpret the metro lines as edges in the drawing of a matching—connecting the respective leaves—in a circular layout. This layout has k bundled crossings, each stemming from a block crossing. Hence, we could use the lower bound of Lemma 2. To this end, we simplify the instance and consider the remaining m lines. Lemma 2 implies that an optimum solution has at least $m/16$ block crossings of the metro lines. We apply the method of Fink et al. [15] creating $2m$ block crossings in $\mathcal{O}(m^2)$ time, and reinsert the simplified lines. \square

6 Conclusion

We have considered the bundled crossing number problem and devised upper and lower bounds for general as well as circular layouts with and without fixed circular vertex order. We have also shown the relation of the bundled crossing number to the orientable graph genus and resolved an open problem for block crossings of metro lines on trees. The setting of bundled crossings still has several interesting questions to offer. It seems very likely that the circular bundled crossing number problem is NP-hard, but a proof is missing. Furthermore, an approximation or a fixed-parameter algorithm for the version with free circular vertex order is desirable. Both questions are also interesting for general graph layouts.

References

1. E. Ackerman and R. Pinchasi. On the degenerate crossing number. *Discrete & Computational Geometry*, 49(3):695–702, 2013.
2. M. Ajtai, V. Chvátal, M. M. Newborn, and E. Szemerédi. Crossing-free subgraphs. *North-Holland Mathematics Studies*, 60:9–12, 1982.
3. M. Baur and U. Brandes. Crossing reduction in circular layouts. In *WG’04*, pages 332–343. Springer, 2004.
4. Q. W. Bouts and B. Speckmann. Clustered edge routing. In *PacificVis’15*, pages 55–62, 2015.
5. C. Buchheim, M. Chimani, C. Gutwenger, M. Jünger, and P. Mutzel. Crossings and planarization. In *Handbook of Graph Drawing and Vis.* CRC Press, 2013.
6. S. Cabello. Hardness of approximation for crossing number. *Discrete & Computational Geometry*, 49(2):348–358, 2013.
7. J. Chuzhoy. An algorithm for the graph crossing number problem. In *STOC’11*, pages 303–312, 2011.
8. W. Cui, H. Zhou, H. Qu, P. C. Wong, and X. Li. Geometry-based edge clustering for graph visualization. *TVCG*, 14(6):1277–1284, 2008.
9. H. de Fraysseix and P. O. de Mendez. Stretching of Jordan arc contact systems. In *GD’03*, volume 2912 of *LNCS*, pages 71–85. Springer, 2003.
10. M. Dickerson, D. Eppstein, M. T. Goodrich, and J. Y. Meng. Confluent drawings: Visualizing non-planar diagrams in a planar way. *JGAA*, 9(1):31–52, 2005.
11. D. Eppstein, D. Holten, M. Löffler, M. Nöllenburg, B. Speckmann, and K. Verbeek. Strict confluent drawing. In *GD’13*, volume 8242, pages 352–363. Springer, 2013.
12. O. Ersoy, C. Hurter, F. V. Paulovich, G. Cantareiro, and A. Telea. Skeleton-based edge bundling for graph visualization. *TVCG*, 17(12):2364–2373, 2011.
13. S. Felsner and P. Valtr. Coding and counting arrangements of pseudolines. *Discrete & Computational Geometry*, 46(3):405–416, 2011.
14. M. Fink, J. Hershberger, S. Suri, and K. Verbeek. Bundled crossings in embedded graphs. In E. Kranakis, G. Navarro, and E. Chávez, editors, *LATIN ’16*, volume 9644 of *LNCS*, pages 454–468. Springer, 2016.
15. M. Fink, S. Pupyrev, and A. Wolff. Ordering metro lines by block crossings. *JGAA*, 19(1):111–153, 2015.
16. E. Gansner, Y. Hu, S. North, and C. Scheidegger. Multilevel agglomerative edge bundling for visualizing large graphs. In *PacificVis’11*, pages 187–194. IEEE, 2011.
17. E. R. Gansner and Y. Koren. Improved circular layouts. In *GD’06*, pages 386–398. Springer, 2007.
18. F. Giordano, G. Liotta, T. Mchedlidze, A. Symvonis, and S. Whitesides. Computing upward topological book embeddings of upward planar digraphs. *J. Discrete Algorithms*, 30:45 – 69, 2015.
19. D. Holten. Hierarchical edge bundles: Visualization of adjacency relations in hierarchical data. *TVCG*, 12(5):741–748, 2006.
20. D. Holten and J. J. van Wijk. Force-directed edge bundling for graph visualization. *Computer Graphics Forum*, 28(3):983–990, 2009.
21. A. Lambert, R. Bourqui, and D. Auber. Winding roads: Routing edges into bundles. *Computer Graphics Forum*, 29(3):853–862, 2010.
22. F. Lazarus, M. Pocchiola, G. Vegter, and A. Verroust. Computing a canonical polygonal schema of an orientable triangulated surface. In *SoCG’01*, pages 80–89. ACM, 2001.
23. B. Mohar. The genus crossing number. *ARS Math. Contemporanea*, 2(2), 2009.

24. J. Pach and G. Tóth. Degenerate crossing numbers. *Discrete & Computational Geometry*, 41(3):376–384, 2009.
25. S. Pupyrev, L. Nachmanson, S. Bereg, and A. E. Holroyd. Edge routing with ordered bundles. *Computational Geometry*, 52:18–33, 2016.
26. M. Schaefer. The graph crossing number and its variants: A survey. *Electronic Journal of Combinatorics*, Dynamic Survey 21, 2013.
27. M. Schaefer and D. Štefankovič. The degenerate crossing number and higher-genus embeddings. In *Graph Drawing and Network Visualization*, pages 63–74. Springer, 2015.
28. F. Shahrokhi, O. Sýkora, L. A. Székely, and I. Vrt'o. Book embeddings and crossing numbers. In *WG'04*, pages 256–268. Springer, 1995.
29. C. Thomassen. The graph genus problem is NP-complete. *Journal of Algorithms*, 10(4):568–576, 1989.
30. K. Yamanaka, S. Nakano, Y. Matsui, R. Uehara, and K. Nakada. Efficient enumeration of all ladder lotteries and its application. *Theor. Comput. Sci.*, 411(16–18):1714–1722, 2010.

Appendix

A Bundled Crossing Number of K_6

Lemma 7. $\text{bc}(K_6) > 1$

Proof. As pointed out in the proof of Lemma 1, we need to verify that there is no embedding of K_6 in the standard square representing the torus in which (i) no two edges incident to the same vertex v leave the square to adjacent sides and (ii) no part of an edge can enter and leave the square on adjacent sides. Hence, for every vertex, edges can leave the square either vertically (to the upper and lower side), horizontally (to the left and right side), or not at all. This yields a partition of V into sets of “vertical” vertices V_v , “horizontal” vertices V_h and free vertices V_f . It is clear that if there is a vertical or horizontal vertex, then there must be at least two vertices of the respective type. However, we know that both sets must be nonempty, since otherwise we would have a planar embedding of K_6 . Hence, $|V_v|, |V_h| \geq 2$.

All edges between free vertices as well as all edges between vertices of different sets (mixed edges) must be drawn entirely inside the square, and, hence, be crossing-free. Since they have edges leaving the square, all vertical and horizontal vertices must lie on the outer face of the planar embedding consisting of the previously mentioned edges not leaving the square. Now, assume that there is at most one free vertex. Then either V_v or V_h contains at least three vertices; since the other set also contains at least two vertices, we find an outerplanar embedding of $K_{2,3}$ in the square. However, $K_{2,3}$ is not outerplanar.

In the remaining case, we must have $|V_v| = |V_h| = |V_f| = 2$. The planar drawing in the sphere consists of a 4-cycle on the outer face connecting the vertical and horizontal vertices and the two free vertices in the interior connected to all other vertices. However, this is not possible in a planar way, a contradiction. Therefore, no simple drawing with only one bundled crossing exists, that is, $\text{bc}(K_6) > 1$. \square

B Triangles in Simplified Circular Instances

Figure 8 shows an example of a circular layout of a simplified instance for which the factor between triangles and edges comes arbitrarily close to $1/4$ for large m . The example is shown with 5 gadgets, but can easily be extended. If we assign each edge to the gadget of its right endvertex, each gadget except for the leftmost and the rightmost one has eight edges. (The leftmost has three and the rightmost has six); each of these gadgets has two triangles (leftmost and rightmost have three). Furthermore, if we extend the example by adding more gadgets to the left, additional edges that are routed through existing gadgets never create additional triangles there since they are routed straight through faces of degree four. Hence,




Fig. 8: Circular layout of a simplified instance with only $m/4 + 3/2$ triangles (shown with 5 gadgets).

for all gadgets except for the leftmost and the rightmost one, the ratio between triangles and edges is $1/4 + 3/(2m)$, which comes arbitrarily close to $1/4$ for large m ; this means that the lower bound of $m/4$ triangles for simplified instances is asymptotically tight.

C Block Crossings in Metro Maps

In this section, we give a short introduction in the setting of minimizing the number of block crossings for metro lines running through a metro network. The input consists of a plane graph $G = (V, E)$, and a set $L = \{l_1, \dots, l_\ell\}$ of simple paths in G . We call G the *underlying network* and the paths (metro) *lines*.

For each edge $e = (u, v) \in E$, let L_e be the set of lines passing through e . For $i \leq j < k$, a *block move* (i, j, k) on the sequence $\pi = [\pi_1, \dots, \pi_n]$ of lines on e is the exchange of two consecutive blocks π_i, \dots, π_j and π_{j+1}, \dots, π_k . Interpreting $e = (u, v)$ directed from u to v , we are interested in *line orders* $\pi^0(e), \dots, \pi^{t(e)}(e)$ on e , so that $\pi^0(e)$ is the order of lines L_e at the beginning of e (that is, at vertex u), $\pi^{t(e)}(e)$ is the order at the end of e (that is, at vertex v), and each $\pi^i(e)$ is an ordering of L_e so that $\pi^{i+1}(e)$ is derived from $\pi^i(e)$ by a block move. If $t+1$ line orders with these properties exist, we say that there are t *block crossings* on edge e .

For block crossings, the *edge crossings* model is used, in which crossings are not hidden under station symbols if possible. Two lines sharing at least one common edge either do not cross or cross each other on an edge but never in a node. Furthermore, we assume the *path terminal property*, that is, any line terminates at a leaf and no two lines terminate at the same leaf; see Fig. 9 for an example. The *block crossing minimization* (BCM) asks for a drawing (given by line orders for the edges) that minimizes the total number of block crossings.




Fig. 9: An example of a BCM instance with 5 lines (taken from [15]).

Note the similarity between block crossings and bundled crossings. If we interpret the metro lines in a drawing as edges of a graph, then every block crossing can be seen as a bundled crossing. However, block crossings are much more restricted compared to bundled crossings: First, lines are restricted to be routed on the metro network and, second, block crossings on an edge have to follow the direction of the edge.

Fink et al. [15] present the following two results that we utilize to prove Theorem 6:

- Assume that l_1 and l_2 are two metro lines for which both pairs of terminals are adjacent in clockwise order of leaves in the given embedding. Then removing one of the lines from the instance does not change the minimum number of block crossings necessary for the instance. This observation corresponds to our simplification step.
- If G is a tree, then there is an algorithm that creates a feasible line ordering with less than 2ℓ block crossings. This result provides an upper bound for a simplified instance.

D Tight Examples for Circular Drawings

Lemma 8. *For every $k \geq 1$, there is a graph $G = (V, E)$ and an outerplanar subgraph $G^* = (V, E)$ with $m^* = m - k$ edges so that any circular vertex order π prescribed by G^* has $\text{bc}^\circ(G, \pi) \geq 2k - 1$.*

Proof. Consider the example shown in Fig. 10. After removing the k red bold edges, the black edges form a maximum outerplanar subgraph; up to rotation a crossing-free drawing of this subgraph results in the vertex order shown in the drawing.

Each bold red edge crosses each other such edge. Furthermore, each bold red edge has crossings with black edges on either side. Two such crossings that don't share the red edge cannot be part of the same bundled crossing. There can be at most one red edge whose black crossings on both ends are in the same bundled crossing; all other red edges have at least two bundled crossings. Hence, there must be at least $2k - 1$ bundled crossings. \square

Lemma 9. *For every $k \geq 1$, there is a graph $G = (V, E)$ with $\text{bc}^\circ(G) = 1$ and maximum outerplanar subgraph $G^* = (V, E)$ with m^* edges so that $m - m^* \geq k$.*




Fig. 10: example

Proof. Figure 11 shows a sketch of the construction of G . With the given circular order of vertices at least one of the two groups of read bold edges (k edges each) must disappear to remove the crossings. On the other hand, even with only the black edges, at least k edges must be removed to destroy all cycles between red vertices; as long as there is such a cycle, it fixes the circular order of these vertices. Hence, at least k edges must be removed to make the graph outerplanar. On the other hand, the drawing clearly shows that a single bundled crossing is sufficient. \square

E Tight Example for General Drawings

Lemma 10. *For every $k \geq 1$, there is a graph $G = (V, E)$ with $\text{bc}(G) = 1$ and maximum planar subgraph $G^* = (V, E)$ with m^* edges so that $m - m^* \geq k$.*

Proof. Figure 12 shows a sketch of the construction of G . The idea is to start with a triangulated graph, choose two (red) vertices far apart, make the triangulation stronger around them and then add k “parallel” edges between neighbors of the red vertices. The graph can only be made planar by removing the k edges, removing a “path” of edges separating the red vertices, or “cutting” one of the red vertices “loose”. All of these involve removing a lot of edges. In contrast to that, as the parallel routing suggests, a single bundled crossing suffices for all k edges. \square




Fig. 11: A graph that can be planarized by removing no less than $k = 3$ edges, yet with genus and bundled crossing number 1.




Fig. 12: A graph that can be planarized by removing no more than k edges, yet with genus and bundled crossing number 1.